arXiv:1701.07926v2 [stat.ML] 13 Feb 2017

Boosting hazard regression with time-varying covariates

Donald K.K. Lee* and Ningyuan Chen

Yale University and HKUST

February 14, 2017

Abstract

Consider a left-truncated right-censored survival process whose evolution depends on time-varying covariates. Given functional data samples from the process, we propose a gradient boosting procedure for estimating its log-intensity function in a flexible manner to capture time-covariate interactions. The estimator is shown to be consistent if the model is correctly specified. Alternatively an oracle inequality can be demonstrated for tree-based models. We use the procedure to shed new light on a question from the operations literature concerning the effect of workload on service rates in an emergency department.

To avoid overfitting, boosting employs several regularization devices. One of them is step-size restriction, but the rationale for this is somewhat mysterious from the viewpoint of consistency: In theoretical treatments of classification and regression problems, unrestricted greedy step-sizes appear to suffice. Given that the partial log-likelihood functional for hazard regression has unbounded curvature, our study suggests that step-size restriction might be a mechanism for preventing the curvature of the risk from derailing convergence.

Keywords: Survival analysis, gradient boosting, time-varying covariates, functional data, step-size restriction, regression trees, partial likelihood, queuing transition intensities, emergency departments.
MSC 2010 subject classifications: 62N02, 62G05, 90B22

1 Introduction

Consider the following left-truncated right-censored (LTRC) survival process: Let U be a random failure time and $X(\cdot) = \{X(t) \in \mathcal{X} \subseteq \mathbb{R}^p\}$ an external, time-varying covariate process. Conditional on surviving up to time $t-$, the probability of failing in $[t, t+dt)$ is $\lambda(t-, X(t-))dt$ where $\lambda(t, x)$ is the intensity function

*Correspondence: Donald Lee (donald.lee@yale.edu), Yale School of Management, New Haven CT 06520, USA

of interest. The survival status of U and the value of $X(t)$ are only observable inside the time window $[L, C]$ defined by the truncation and censoring times respectively, with censoring occurring no later than a fixed time that is normalized to 1. Thus given N realizations $\{L_i, U_i, C_i, X_i(\cdot)\}_{i=1}^N$ of the process, those for which failure occurred before truncation ($U_i < L_i$) are unobserved. The remaining $n \leq N$ realizations generate the samples $\{L_i, T_i, \Delta_i, X_i(t)_{t\in[L_i,T_i]}\}_{i=1}^n$ where $T_i = \min(U_i, C_i)$ and $\Delta_i = I(L_i \leq U_i \leq C_i)$ is the censoring indicator. It is assumed that U and the truncation-censoring pair (L, C) are conditionally independent given $X(\cdot)$.

In addition to being of vital importance to the health sciences, the process above also underpins a wide swath of other applications ranging from credit default models to customer transition dynamics in a queuing network. Having accurate models for these enables one to build high resolution simulations and/or to evaluate the goodness-of-fit of simpler parametric forms. To this end, we wish to estimate the log-intensity $F(t, x) = \log \lambda(t, x)$ from a given set of functional data samples $\{L_i, T_i, \Delta_i, X_i(t)_{t\in[L_i,T_i]}\}_{i=1}^n$. To capture possible interactions in the time-covariate domain of $F(t, x)$, a flexible estimator is desired. A few nonparametric approaches exist in literature and are described at the end of this section. Some of them are suited to scenarios involving a small number of covariates, while others focus on the analysis of theoretical estimators. In this paper we propose a practical boosting procedure for estimating $F(t, x)$ in some class of functions defined on the time-covariate domain $[0, 1] \times \mathcal{X}$. Gradient boosting (Breiman [7], Friedman [15], Mason et al. [39]) is a regularized functional descent algorithm for minimizing risk, which in this case is the negative partial log-likelihood[1]

$$\hat{R}_n(F) = \frac{1}{n} \sum_{i=1}^{n} \left\{ \int_{L_i}^{T_i} e^{F(t,X_i(t))} dt - \Delta_i F(T_i, X_i(T_i)) \right\}. \tag{1.1}$$

Given that $\hat{R}_n(F)$ can be unbounded below, which can in turn lead to overfitting, three types of regularization are employed: i) A weak learner approximation to the gradient function is used as the descent direction (e.g. coordinate descent or a shallow regression tree fit [15]) which can also induce variable selection; ii) A small step-size is used for the line update [15]; and iii) The number of descent iterations is capped. Algorithm 1 describes the generic gradient boosting procedure, while our proposed procedure for hazard regression will be presented in section 2 (Algorithm 2).

In contrast to common classification and regression problems that involve non-functional data, the way that $\hat{R}_n(F)$ depends on F poses a significant challenge to implementing boosting for hazard regression: We will show that the gradient of $\hat{R}_n(F)$ cannot be computed from its canonical representation (1.1). The key to rectifying this is to carefully specify the domain for $\hat{R}_n(F)$. In section 2 we construct a sample-dependent domain on which (1.1) can be re-expressed

[1]The term $F(T_i, X_i(T_i))$ in the likelihood should strictly speaking be $F(T_i-, X_i(T_i-))$, however for our purposes it suffice to work with the modification (1.1) as its expectation is also uniquely minimized at $\log \lambda$ (see Lemma 2).

Algorithm 1 Generic gradient boosting

1 Initialize $F_0(t, x)$.

2 For $m = 0$ to $\hat{m}$:

- Compute a simple approximation $\tilde{g}_m(t, x)$ to the gradient of $\hat{R}_n(F)$ at F_m.
- Perform the updates $F_{m+1}(t, x) \leftarrow F_m(t, x) - \theta_m \tilde{g}_m(t, x)$, $m \leftarrow m + 1$ using a small step-size θ_m.

3 The boosting estimator is $\hat{F}_{\hat{m}}(t, x) = F_0(t, x) - \sum_{m=0}^{\hat{m}-1} \theta_m \tilde{g}_m(t, x)$.

as an integral with respect to a random measure $\hat{\mu}_n$ on $[0, 1] \times \mathcal{X}$. Specifically, given a linear space $\mathcal{F}$ of functions on $[0, 1] \times \mathcal{X}$ we use its completion in $L^2(\hat{\mu}_n)$, $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\hat{\mu}_n})$, as the domain. The tradeoff is that, unlike (1.1), the integral representation is not coordinate-free and this has implications for computations. To ensure that $\hat{R}_n(F) > -\infty$ on $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\hat{\mu}_n})$, we require each function in $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\hat{\mu}_n})$ to be bounded on the support of $\hat{\mu}_n$. Note that $\mathcal{F}$ must then be finite-dimensional because the resulting inclusion $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\hat{\mu}_n}) \hookrightarrow L^\infty(\hat{\mu}_n)$ is an embedding. Under regularity conditions our estimator is consistent if $\mathcal{F}$ is correctly specified. On the other hand if $\mathcal{F}$ comprises of flexible piecewise constant functions like regression trees, then the plug-in intensity estimator satisfies an oracle inequality.

The number of boosting iterations and the degree of step-size restriction employed in our procedure will be justified in section 3. The former is set using the framework of Zhang and Yu [58], which shows for general loss functions that early stopping serves to prevent the estimator from growing too large too fast. On the other hand the role of step-size restriction is more mysterious. While [58] demonstrates that small step-sizes are needed to minimize empirical risk in order to prove consistency, it appears that unrestricted greedy step-sizes are already small enough for commonly encountered learning problems: For the usual classification losses, Telgarsky [52] and the works cited therein show that shrinkage is only needed for margin maximization. For other losses that are strongly convex and have uniformly bounded curvatures, the appendix of [58] shows that greedy line-search suffices for minimizing risk. In our study of hazard regression where the empirical risk has unbounded curvature, we contribute to clarifying when explicit step-size restrictions are necessary: We find that it serves as a mechanism for counterbalancing the curvature of the risk. If the curvature is unbounded then the step-sizes may need to be explicitly controlled to ensure empirical risk minimization. Our parsimonious bounds make it easy to see that smaller step-sizes always lead to faster convergence (in terms of sample size) at the expense of more iterations, with diminishing returns in performance. This matches the empirical observations in [15] and holds regardless of whether the risk has bounded curvature or not.

The remainder of this paper is organized as follows. Related approaches in literature will be described below. The proposed estimator will be constructed in section 2 and its guarantees derived in section 3. In section 4 we apply our procedure to a dataset from an emergency department to shed new light on a question from the operations literature concerning the effect of workload on service rates. Concluding remarks can be found in section 5.

Related nonparametric approaches. Since $\hat{R}_n(F)$ can be unbounded below, some restrictions need to be imposed on the function class $\mathcal{F}$ or on $F(t, x)$ itself to prevent overfitting. One way is to impose shape constraints on the intensity function, as first suggested in Grenander [17]. Although there are methods for shape-constrained multivariate density estimation like Cule et al. [12], no analogue appears to exist for hazard regression in the presence of covariates.

Another way to obtain restrictions is to control the dimension of $\mathcal{F}$, allowing it to grow slowly with sample size. The resulting sequence of MLEs form a sieve estimator that can handle time-fixed covariates (Dohler and Ruschendorf [14], Karr [24], Kooperberg et al. [30]) as well as time-varying ones (Huang and Stone [23]). However, for a given sample size, the theory provides no procedure for computing the functional MLE or even which $\mathcal{F}$ to use. The MLE is particularly challenging to compute when additional constraints are assumed as in [14]. In practice, researchers resort to using a stepwise feature selection heuristic[2] to construct an estimator from products of univariate splines (Kooperberg and Clarkson [28], Kooperberg et al. [29]).

An opposite approach to sieves is to fix a large or flexible function class and directly control the complexity of the estimator itself. Both boosting and penalized likelihood methods are similar in this sense. However, beyond the case of a single time-fixed covariate (Gu [18], O'Sullivan [44], Senthilselvan [50]), recent literature on penalized likelihood have moved from the nonparametric setting toward some form of separability.

For methods not based on likelihood maximization, kernel estimators can be obtained from smoothing in time-covariate space (Li and Doss [32], McKeague and Utikal [41], Nielsen and Linton [43], Perez et al. [45], Spierdijk [51], van Keilegom and Veraverbeke [56]). If the covariates are high-dimensional then kernel smoothing is susceptible to the curse of dimensionality, so semi-parametric restrictions are needed (Linton et al. [34]).

Boosting semiparametric hazard models. Gradient boosting have previously been applied to both parametric and semiparametric models with time-fixed covariates: Ridgeway [46], Li and Luan [33], and Hothorn et al. [22] considered the Cox proportional hazards model while Schmid and Hothorn [49] examined accelerated failure time models. For additive hazards models, high-dimensional covariates can be handled using either principal components re-

[2] An alternative is to cast sieve estimation as a model selection problem where $\hat{F}$ and $\dim \mathcal{F}$ are jointly chosen to minimize the sum of the risk and a penalty on $\dim \mathcal{F}$. However this requires solving a separate minimization problem for each candidate function class, which is computationally burdensome except for cases where the sequence of minimizers have orthogonal increments.

gression (Ma et al. [37]) or Lasso (Martinussen and Scheike [38]). Boosting can certainly be applied as well given that the gradient approximation step in Algorithm 1 can be used to select variables.

2 Constructing the estimator

The chief difficulty with implementing gradient boosting for hazard regression is that the gradient of the likelihood risk (1.1) cannot be determined from its directional derivative

$$\frac{d}{d\theta}\hat{R}_n(F+\theta f)\bigg|_{\theta=0} = \frac{1}{n}\sum_{i=1}^{n}\int_{L_i}^{T_i} e^{F(t,X_i(t))}f(t,X_i(t))dt - \frac{1}{n}\sum_{i=1}^{n}\Delta_i f(T_i,X_i(T_i)).$$

While this can be written as the difference of two different inner products $\langle e^F, f\rangle_\dagger - \langle 1, f\rangle_\ddagger$, these cannot be combined into a single inner product of f with some gradient function. This is intuitive since the first one involves time integrals while the second consists of pointwise evaluations. To remedy this, we will equip the class $\mathcal{F}$ of candidate estimators with an appropriate reproducing kernel by carefully specifying a sample-dependent domain for $\hat{R}_n(F)$. The likelihood risk can then be re-expressed as a smooth convex functional whose gradient is readily obtainable. An analogous representation also exists for the expected risk which forms the basis for analyzing the estimator.

Laying the groundwork. We begin by stating four conditions A1-A4 that will be used throughout the paper. Condition A1 below is the same as Condition 1(iv) of [23].

A1 $\lambda(t,x)$ is bounded between some interval $[\Lambda_L, \Lambda_H] \subset (0,\infty)$ on the time-covariate domain $[0,1]\times\mathcal{X}$.

A2 $\{L_i, U_i, C_i, X_i(\cdot)\}_{i=1}^N$ are independent and identically distributed[3].

Next, some regularity conditions on the covariate process $X(\cdot)$ are needed for the existence of the expected likelihood risk (Lemma 2). Let $\mathcal{T}$ denote the Borel algebra for $[0,1]$, and $(\Omega,\mathcal{S},\mathbb{P})$ the probability space for L, U, C, and $X(\cdot)$. Implicit in this statement is that $X(\cdot)$ is a random function, which we require to satisfy

A3 The map $(\omega,t)\mapsto X(t;\omega)$ is $\mathcal{S}\otimes\mathcal{T}$-measurable.

The condition ensures that $\{t, X(t)\}$ is $\mathcal{S}\otimes\mathcal{T}$-measurable and $\{T, X(T)\}$ is $\mathcal{S}$-measurable, and it is automatically satisfied if the sample trajectories of $X(t)$ are either left-continuous or right-continuous (progressive measurability). Denoting the indicator function as $I(\cdot)$, we can then define the following population and empirical sub-probability measures on $[0,1]\times\mathcal{X}$:

$$\mu(B) = \mathbb{E}\int_0^1 I(L\le t\le T)I[\{t,X(t)\}\in B]dt, \tag{2.1}$$

[3]In the context of the application in section 4, simulation results suggest that our procedure is potentially robust to non-IID covariates.

$$\begin{aligned}\hat{\mu}_n(B) &= \frac{1}{n}\sum_{i=1}^{n}\int_{L_i}^{T_i} I[\{t, X_i(t)\} \in B]dt \\ &= \frac{1}{n}\sum_{i=1}^{N}\int_{0}^{1} I(L_i \leq t \leq T_i) I[\{t, X_i(t)\} \in B]dt.\end{aligned} \tag{2.2}$$

Intuitively $\hat{\mu}_n$ measures the denseness of the sample time-covariate paths on $[0,1]\times\mathcal{X}$. Bearing in mind that $n = \sum_{i=1}^{N} I(L_i \leq U_i)$ is random, $\mathbb{E}\{(n/N)\hat{\mu}_n(B)\} = \mu(B)$ under A2. Furthermore for any integrable f,

$$\int f d\mu = \mathbb{E}\int_0^1 I(L \leq t \leq T) f(t, X(t))dt, \tag{2.3}$$

$$\int f d\hat{\mu}_n = \frac{1}{n}\sum_{i=1}^{n}\int_{L_i}^{T_i} f(t, X_i(t))dt. \tag{2.4}$$

This allows us to define the following (random) norms and inner product

$$\begin{aligned}\|f\|_{\hat{\mu}_n,1} &= \int |f| d\hat{\mu}_n, \\ \|f\|_{\hat{\mu}_n,2} &= \left(\int f^2 d\hat{\mu}_n\right)^{1/2}, \\ \|f\|_{\infty} &= \sup\left\{|f(t,x)| : (t,x) \in [0,1]\times\mathcal{X}\right\}, \\ \langle f, f'\rangle_{\hat{\mu}_n} &= \int f f' d\hat{\mu}_n,\end{aligned}$$

and note that $\|\cdot\|_{\hat{\mu}_n,1} \leq \|\cdot\|_{\hat{\mu}_n,2} \leq \|\cdot\|_{\infty}$ because $\hat{\mu}_n([0,1]\times\mathcal{X}) \leq 1$.

By design, $\hat{\mu}_n$ allows us to specify a natural domain for $\hat{R}_n(F)$. Let $\{\phi_j\}_{j=1}^{d}$ be a set of bounded normalized functions $[0,1]\times\mathcal{X} \mapsto [-1,1]$ that are linearly independent, in the sense[4] that $\int_{[0,1]\times\mathcal{X}}(\sum_j c_j\phi_j)^2 dtdx = 0$ if and only if $c_1 = \cdots = c_d = 0$. The span of the functions is

$$\mathcal{F} = \left\{\sum_{j=1}^{d} c_j\phi_j : c_j \in \mathbb{R}\right\}.$$

When $\mathcal{F}$ is equipped with $\langle\cdot,\cdot\rangle_{\hat{\mu}_n}$ we obtain the following sample-dependent subspace of $L^2(\hat{\mu}_n)$:

$$(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n}). \tag{2.5}$$

Since $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ comprises of equivalence classes $[F] = \{F' \in \mathcal{F} : F' = F\ \hat{\mu}_n\text{-a.e.}\}$ rather than actual functions, $\dim(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n}) \leq d$. Furthermore the pointwise values of $[F]$ need to be defined in order for the likelihood risk (1.1) to be defined on $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$. Henceforth, given a basis[5] $\{\hat{\varphi}_j(t,x)\}_j \subset \mathcal{F}$ for

[4] If some of the covariates are discrete-valued then dx should be interpreted as the product of a counting measure and a Lebesgue measure.

[5] Computed from applying Gram-Schmidt to $\phi_1(t,x), \cdots, \phi_d(t,x)$ for example.

$(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$, we shall represent each class with the unique function of the form $\sum_j c_j \hat{\varphi}_j(t,x)$ contained within it. However, the resulting definition of $\hat{R}_n([F])$ is not coordinate-free: Two different bases may lead to different pointwise values[6] for $[F]$ and hence different values for $\hat{R}_n([F])$. Nonetheless the guarantees for our estimator in section 3 hold regardless of basis choice.

A simple example will help illustrate the discussion above. Consider the case of no covariates and only one observation, a failure at $T = 1/4$. Suppose $\mathcal{F}$ is spanned by

$$\{\phi_1(t) = 1, \phi_2(t) = I_{[0,1/4)}(t)\}. \tag{2.6}$$

Then $\int f d\hat{\mu}_n = \int_0^{1/4} f(t)dt$, and two possible orthonormal bases for $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ are $\{\hat{\varphi}_1(t) = 2\phi_1(t)\}$ and $\{\hat{\psi}_1(t) = 2\phi_2(t)\}$. In this case $\dim(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n}) < \dim \mathcal{F}$ because ϕ_1 and ϕ_2 belong to the same equivalence class in $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$. Moreover $\hat{\varphi}_1(t)$ and $\hat{\psi}_1(t)$ differ at the point $t = 1/4$ even though they also belong to the same class.

The final condition we impose is for $\{\phi_j\}_{j=1}^d$ to be linearly independent in $L^2(\mu)$, that is $\|\sum_j c_j\phi_j\|_{\mu,2}^2 = \sum_{ij} c_i \left(\int \phi_i\phi_j d\mu\right) c_j = 0$ if and only if $c_1 = \cdots = c_d = 0$. Since by construction $\{\phi_j\}_{j=1}^d$ is already linearly independent on $[0,1]\times\mathcal{X}$, the condition intuitively requires the set of all possible time-covariate trajectories to be adequately dense in $[0,1]\times\mathcal{X}$ to intersect a sufficient amount of the support of every ϕ_j. This is weaker than conditions 1(ii)-1(iii) in [23] which require $X(t)$ to have a positive joint probability density on $[0,1]\times\mathcal{X}$, thereby ruling out discrete and categorical covariates.

A4 The Gram matrix $\Sigma_{ij} = \int \phi_i\phi_j d\mu$ is positive definite.

Since $\mathcal{F}$ is finite-dimensional, all norms on it are equivalent so the following is finite:

$$\alpha_{\mathcal{F}} = 2 \sup_{F\in\mathcal{F}:\|F\|_\infty=1} \left(\frac{\|F\|_\infty}{\|F\|_{\mu,1}}\right) = 2/\inf_{F\in\mathcal{F}:\|F\|_\infty=1} \|F\|_{\mu,1} > 1, \tag{2.7}$$

and for all $F \in \mathcal{F}$

$$\|F\|_{\mu,1} \le \|F\|_{\mu,2} \le \|F\|_\infty \le \frac{\alpha_{\mathcal{F}}}{2}\|F\|_{\mu,1}. \tag{2.8}$$

The factor of 2 in the definition of (2.7) is arbitrary and can be replaced with anything greater than 1. It will be shown in Lemma 3 that with high probability there exists an empirical analogue to (2.8).

Integral representations for the likelihood risk. By design $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ admits a reproducing kernel, which can be used to recast $\hat{R}_n(F)$ as a smooth convex functional on $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ with gradient (2.11). As discussed, the representation is not coordinate-free and will depend explicitly on the chosen basis $\{\hat{\varphi}_j(t,x)\}_j$. Hence it is only valid for functions in the pointwise span of the basis.

[6] Unless $\dim(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n}) = d$, in which case each class contains exactly one function.

Lemma 1. *With respect to an orthonormal basis $\{\hat{\varphi}_j(t,x)\}_j$ of $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$, for $F, f \in$ span$\{\hat{\varphi}_j(t,x)\}_j$ the likelihood risk (1.1) can be written as*

$$\hat{R}_n(F) = \int (e^F - \hat{\lambda}F) d\hat{\mu}_n, \tag{2.9}$$

where $\hat{\lambda} \in (\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ is the function

$$\hat{\lambda}(t,x) = \frac{1}{n} \sum_j \left\{ \sum_{i=1}^n \Delta_i \hat{\varphi}_j(T_i, X_i(T_i)) \right\} \hat{\varphi}_j(t,x).$$

Thus there exists $\hat{\rho} \in (0,1)$ (depending on F and f) for which the Taylor representation

$$\hat{R}_n(F+f) = \hat{R}_n(F) + \langle \hat{g}_F, f \rangle_{\hat{\mu}_n} + \frac{1}{2} \int e^{F+\hat{\rho}f} f^2 d\hat{\mu}_n \tag{2.10}$$

holds, where the gradient

$$\hat{g}_F(t,x) = \sum_j \left\langle e^F, \hat{\varphi}_j \right\rangle_{\hat{\mu}_n} \hat{\varphi}_j(t,x) - \hat{\lambda}(t,x) \tag{2.11}$$

of $\hat{R}_n(F)$ is the projection of $e^F - \hat{\lambda}$ onto $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$. Hence if $\hat{g}_F = 0$ then the infimum of $\hat{R}_n(F)$ over span$\{\hat{\varphi}_j(t,x)\}_j$ is uniquely attained at F.

Remark. To see what can go wrong when (2.9) is applied to functions outside span$\{\hat{\varphi}_j(t,x)\}_j$, consider the example from (2.6). If the basis $\{\hat{\varphi}_1(t) = 2\phi_1(t)\}$ is chosen, then $\phi_2(t)$ is not pointwise spanned by $\hat{\varphi}_1(t)$ even though ϕ_2 and $\hat{\varphi}_1/2$ belong to the same equivalence class. Under (1.1) we have $\hat{R}_n(\phi_2) = e/4$ while $\int (e^{\phi_2} - \hat{\lambda}\phi_2) d\hat{\mu}_n = e/4 - 1$.

Remark. The infimum of $\hat{R}_n(F)$ is not always attainable, so the log-intensity MLE might not even exist: If f is non-positive and vanishes on the set $\{\{T_i, X_i(T_i)\} : \Delta_i = 1\}$, then $\hat{R}_n(F + \theta f) = \int (e^{F+\theta f} - \hat{\lambda}F) d\hat{\mu}_n$ is decreasing in θ so f is a direction of recession. This is however not an issue for our estimator because of early stopping.

The expectation of the (scaled) likelihood risk also has an integral representation which forms the basis for the theory in section 3. A special case of the representation (2.12) below was derived in equation 4.1 of [23], under assumptions more stringent than A4 that do not allow for discrete or categorical covariates. In our setting, the key to establishing (2.12) is to show that on $[0,1] \times \mathcal{X}$, the point process $M(B) = \Delta \cdot I[\{T, X(T)\} \in B]$ (for the point of failure, if observed) has mean $\int_B \lambda d\mu$, from which we deduce Campbell's formula $\mathbb{E}\{\Delta \cdot F(T, X(T))\} = \int F\lambda d\mu$.

Lemma 2. *For $F \in \mathcal{F} \cup \{\log \lambda\}$,*

$$R(F) = \mathbb{E}\left\{ \frac{n}{N} \hat{R}_n(F) \right\} = \int (e^F - \lambda F) d\mu. \tag{2.12}$$

Furthermore the restriction of $R(F)$ to $\mathcal{F}$ is coercive:

$$\frac{1}{2}R(F) \geq \frac{\Lambda_L}{\alpha_{\mathcal{F}}}\|F\|_\infty + \Lambda_H \min\{0, 1 - \log(2\Lambda_H)\}, \tag{2.13}$$

and it attains its minimum at a unique point F^*. If $\mathcal{F}$ contains the underlying log-intensity function then $F^* = \log\lambda$.

Remark. Coerciveness (2.13) implies that any F with expected risk $R(F)$ less than $R(0) \leq 1 < 3$ is uniformly bounded:

$$\begin{aligned} \|F\|_\infty &< \frac{\alpha_{\mathcal{F}}}{\Lambda_L}[3/2 + \Lambda_H \max\{0, \log(2\Lambda_H) - 1\}] \\ &\leq \alpha_{\mathcal{F}}\beta_\Lambda \end{aligned} \tag{2.14}$$

where the constant

$$\beta_\Lambda = \frac{3/2 + \Lambda_H \max\{0, \log(2\Lambda_H) - 1\}}{\min\{1, \Lambda_L\}} \tag{2.15}$$

is by design no smaller than 1 in order to simplify subsequent analyses.

The boosting procedure. In the generic Algorithm 1, three types of regularization mechanisms are employed. First, a weak learner approximation to the negative gradient $-\hat{g}_F$ (2.11) is used as the descent direction. A few choices that induce some degree of variable selection are:

- Small (limited depth) regression trees that are correlated with $-\hat{g}_F$. This provides some smoothing of the gradient function [15].
- The member of the basis $\{\phi_j\}_j$ most aligned with $-\hat{g}_F$. This is a variant of coordinate descent, see for example Chapter 7.2 of Schapire and Freund [48].
- If coordinate descent is used and each ϕ_j depends on only one component of the covariate vector, we obtain componentwise learners (Bühlmann and Yu [11]). The resulting estimator is additive in the covariate components.

To model a generic approximation to a non-zero gradient, we say that a unit vector $\hat{g}_F^\epsilon \in (\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ is an ϵ-gradient at F if for some $0 \leq \epsilon < 1$,

$$\left\langle \frac{\hat{g}_F}{\|\hat{g}_F\|_{\hat{\mu}_n,2}}, \hat{g}_F^\epsilon \right\rangle_{\hat{\mu}_n} \geq (1 - \epsilon). \tag{2.16}$$

The smaller ϵ is, the closer the alignment is between the gradient and the ϵ-gradient. In particular $\hat{g}_F$ is the unique 0-gradient. On the other hand, very simple descent directions will be less aligned with $\hat{g}_F$, resulting in a larger ϵ.

The other two types of regularizations are specified below for our procedure. These choices serve to simplify the results in section 3:

Algorithm 2 Boosting algorithm for hazard regression

1 Set $\hat{F}_0 = 0$, $m = 0$ and repeat the following steps:

- Determine the current gradient $\hat{g}_{\hat{F}_m}$ (2.11). If $\hat{g}_{\hat{F}_m} = 0$ then exit loop, else compute a simple approximation $\hat{g}^{\epsilon}_{\hat{F}_m}$ that satisfies (2.16).
- If
$$\left\| \hat{F}_m - \frac{\nu_n}{m+1} \hat{g}^{\epsilon}_{\hat{F}_m} \right\|_{\infty} < \Psi_n \tag{2.19}$$
then update $\hat{F}_{m+1} \leftarrow \hat{F}_m - \frac{\nu_n}{m+1} \hat{g}^{\epsilon}_{\hat{F}_m}$, $m \leftarrow m + 1$. Else exit loop.

2 Set $\hat{m} \leftarrow m$. The estimators for the log-intensity and intensity functions are respectively
$$\hat{F}_{\hat{m}} = -\sum_{m=0}^{\hat{m}-1} \frac{\nu_n}{m+1} \hat{g}^{\epsilon}_{\hat{F}_m}, \ \hat{\lambda}_{boost} = e^{\hat{F}_{\hat{m}}}.$$

- The number of boosting iterations $\hat{m}$ is controlled by stopping the algorithm before the uniform norm of the estimator $\|\hat{F}_{\hat{m}}\|_{\infty}$ reaches or exceeds
$$\Psi_n = W(n^{1/4}) \to \infty, \tag{2.17}$$
where $W(y)$ is the branch of the Lambert function that returns the real root of the equation $ze^z = y$ for $y > 0$.
- The step-size for the m-th iteration is set to $\nu_n/(m+1)$, which is controlled in two ways. First, it is made to decrease with each iteration according to the Robbins-Monro [47] condition that the sum of the steps diverges while the sum of squared steps converges. The specific rate $1/(m+1)$ above is chosen out of mere convenience. Second, the shrinkage factor ν_n makes the step-sizes decay with n at rate
$$\nu_n^2 e^{\Psi_n} < 1, \ \nu_n^2 e^{\Psi_n} \to 0. \tag{2.18}$$
This acts as a counterbalance to $\hat{R}_n(F)$'s curvature: $\frac{d^2}{d\theta^2} \hat{R}_n(F + \theta f)\Big|_{\theta=0} = \int e^F f^2 d\hat{\mu}_n \leq e^{\Psi_n}$ if $\|F\|_{\infty} < \Psi_n$ and $\|f\|_{\hat{\mu}_n,2} = 1$.

Algorithm 2 describes the proposed boosting procedure for estimating $\log \lambda$.

3 Properties

Guarantees for our estimator can be derived for two scenarios of particular interest. In Proposition 5 we first consider models that are correctly specified

($\mathcal{F}$ includes the true log-intensity function). As such, the minimizer of $R(F)$ is $F^* = \log\lambda$ (Lemma 2) and $\hat{F}_{\hat{m}}$ is consistent:

$$\left\|\hat{F}_{\hat{m}} - \log\lambda\right\|^2_{\mu,2} < \text{error bound.}$$

Proposition 6 concerns piecewise constant functions like regression trees. A key feature of such classes is that they are closed under pointwise exponentiation ($e^F \in \mathcal{F}$ for all $F \in \mathcal{F}$). As a result it can be shown that among all candidate intensity estimators $\{e^F : F \in \mathcal{F}\}$, $\lambda^* = e^{F^*}$ is the one closest to λ in $L^2(\mu)$. An oracle inequality for $\hat{\lambda}_{boost} \to \lambda^*$ can then be demonstrated:

$$\left\|\hat{\lambda}_{boost} - \lambda\right\|^2_{\mu,2} = \min_{F\in\mathcal{F}} \left\|e^F - \lambda\right\|^2_{\mu,2} + \text{error bound.}$$

The value of λ^* over one of its piecewise constant regions B is given by the weighted average $\int_B \lambda d\mu/\mu(B)$ of $\lambda(t,x)$ over B. Hence for sufficiently fine partitions of the time-covariate domain $[0,1]\times\mathcal{X}$, λ^* will closely approximate the value of λ in the interior of the support of μ.

The representation of $R(F)$ and its coercivity in Lemma 2 provide the basis for formalizing Propositions 5 and 6. First, applying Taylor's theorem to (2.12) about F^* yields

$$\left\|\hat{F}_{\hat{m}} - F^*\right\|^2_{\mu,2} \le 2\frac{R(\hat{F}_{\hat{m}}) - R(F^*)}{\min_{t,x}(\lambda^* \wedge \hat{\lambda}_{boost})}. \tag{3.1}$$

The problem is thus transformed into one of risk minimization $R(\hat{F}_{\hat{m}}) \to R(F^*)$, for which [58] suggested analyzing separately the terms of the decomposition

$$\begin{aligned} 0 \le R(\hat{F}_{\hat{m}}) - R(F^*) < &\left|\frac{n}{N}\hat{R}_n(\hat{F}_{\hat{m}}) - R(\hat{F}_{\hat{m}})\right| \\ &+ \left|\frac{n}{N}\hat{R}_n(F^*) - R(F^*)\right| \\ &+ \frac{n}{N}\{\hat{R}_n(\hat{F}_{\hat{m}}) - \hat{R}_n(F^*)\}. \end{aligned} \tag{3.2}$$

The authors argue that in boosting, the point of limiting the number of iterations $\hat{m}$ (enforced by (2.19) in our procedure) is to prevent $\hat{F}_{\hat{m}}$ from growing too fast, so that the first RHS term in (3.2) converges to zero as $n \to \infty$. At the same time $\hat{m}$ is allowed to grow with n, so the empirical risk $\hat{R}_n(\hat{F}_{\hat{m}})$ in the last term of (3.2) is eventually minimized as $n \to \infty$. Lemmas 3 and 4 below show that our procedure achieves both goals. Note that the second RHS term in (3.2) can be bounded using standard concentration results.

First RHS term in (3.2). In the study of sieve hazard estimation for time-fixed covariates, [14] provides a maximal inequality for bounding a similar term. Rather than extend that result to our functional data setting, one that is more suited to our purposes is given in Lemma 3: To capture the qualitative effect of using a simple ϵ-gradient (2.16) as the descent direction, we bound the term in terms of the complexity of

$$\mathcal{F}_\epsilon \subseteq \mathcal{F}, \tag{3.3}$$

the collection of all possible iterates $\hat{F}_m = -\sum_{k=0}^{m-1} \frac{\nu_n}{k+1} \hat{g}^{\epsilon}_{\hat{F}_k}$ for $m = 0, 1, \cdots$. For technical convenience we further enlarge $\mathcal{F}_\epsilon$ to include the unit ball $\{F \in \mathcal{F} : \|F\|_\infty = 1\}$. Depending on the choice of weak learners for the ϵ-gradients, $\mathcal{F}_\epsilon$ may be much less 'complex' than $\mathcal{F}$. For example, coordinate descent might only ever select a small subset of basis functions $\{\phi_j\}_j$ because of sparsity.

The measure of complexity we use below comes from empirical process theory: Define $\mathcal{F}^\Psi_\epsilon = \{F \in \mathcal{F}_\epsilon : \|F\|_\infty < \Psi\}$ and suppose that Q is a sub-probability measure on $[0,1] \times \mathcal{X}$. Then the $L^2(Q)$-ball of radius ε centered at some $F \in L^2(Q)$ is $\left\{F' \in \mathcal{F}^\Psi_\epsilon : \|F' - F\|_{Q,2} < \varepsilon\right\}$. The covering number $\mathcal{N}(\varepsilon, \mathcal{F}^\Psi_\epsilon, Q)$ is the minimum number of such balls needed to cover $\mathcal{F}^\Psi_\epsilon$ (Definitions 2.1.5 and 2.2.3 of van der Vaart and Wellner [55]), so $\mathcal{N}(\varepsilon, \mathcal{F}^\Psi_\epsilon, Q) = 1$ for $\varepsilon \geq \Psi$. A complexity measure for $\mathcal{F}_\epsilon$ is

$$J_{\mathcal{F}_\epsilon} = \sup_{\Psi, Q} \int_0^1 \{\log \mathcal{N}(u\Psi, \mathcal{F}^\Psi_\epsilon, Q)\}^{1/2} du, \tag{3.4}$$

where the supremum is taken over $\Psi > 0$ and over all non-zero sub-probability measures.

As discussed, $J_{\mathcal{F}_\epsilon}$ can potentially be much smaller than $J_{\mathcal{F}}$, which is itself no greater than a universal multiple of $d^{1/2}$ (Theorem 2.6.7 of [55]). However, as pointed out after (2.16), overly simple weak learners may result in a poorer approximation of the gradient and hence a larger value of ϵ. The bounds in Propositions 5 and 6 capture the way this tension induces a performance trade-off.

Lemma 3. *There exists a universal constant κ such that for any $0 < \eta < 1$, with probability at least*

$$1 - 4\exp\left\{-\left(\frac{\eta n^{1/4}}{\kappa \alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon}}\right)^2\right\}$$

an empirical analogue to the norm equivalences (2.8) holds for all $F \in \mathcal{F}$:

$$\|F\|_{\hat{\mu}_n,1} \leq \|F\|_{\hat{\mu}_n,2} \leq \|F\|_\infty \leq \alpha_{\mathcal{F}} \|F\|_{\hat{\mu}_n,1}, \tag{3.5}$$

and for all $F \in \mathcal{F}^{\Psi_n}_\epsilon$,

$$\left|\frac{n}{N}\{\hat{R}_n(F) - \hat{R}_n(0)\} - \{R(F) - R(0)\}\right| < \eta. \tag{3.6}$$

Remark. The equivalences (3.5) imply that $\dim(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n}) = d$: If $\|\sum_j c_j \phi_j\|_{\hat{\mu}_n,2} = 0$ then $\|\sum_j c_j \phi_j\|_\infty = 0$, so $c_1 = \cdots = c_d = 0$ because $\{\phi_j\}_{j=1}^d$ are linearly independent on $[0,1] \times \mathcal{X}$.

Last term in (3.2). We use the representation (2.9) to study the minimization of the empirical risk $\hat{R}_n(F)$ by boosting. Standard results for exact gradient descent like Theorem 2.1.15 of Nesterov [42] are in terms of the norm of the minimizer, which may not exist for $\hat{R}_n(F)$ (see remark after Lemma 1).

If coordinate descent is used instead, section 4.1 of [58] can be applied to convex functions whose infimum may not be attainable, but its curvature is required to be uniformly bounded above. Since the second derivative of $\hat{R}_n(F)$ is unbounded, Lemma 4 below provides two remedies: i) Use the shrinkage decay (2.18) of ν_n to counterbalance the curvature; ii) Use coercivity (2.14) to show that with increasing probability, the iterates $\hat{F}_m$ are uniformly bounded, so the curvatures at those points are also uniformly bounded. Lemma 4 combines both to derive a result that is simpler than what can be achieved from either one alone. It relies in part on adapting the analysis in Lemma 4.1 of [58] for coordinate descent to the case for generic ϵ-gradients. The conditions required below will be shown to hold with high probability in Proposition 5.

Lemma 4. *Suppose (3.5) holds and that*

$$\left|\frac{n}{N}\hat{R}_n(F^*) - R(F^*)\right| < 1, \sup_{F \in \mathcal{F}_\epsilon^{\Psi_n}} \left|\frac{n}{N}\hat{R}_n(F) - R(F)\right| < 1.$$

Then the largest gap between F^* and the iterates $\hat{F}_m$,

$$\hat{\gamma} = \max_{m \leq \hat{m}} \|\hat{F}_m - F^*\|_\infty \vee 1, \tag{3.7}$$

is bounded by a constant no greater than $2\alpha_{\mathcal{F}}\beta_\Lambda$, and for $n \geq 55$,

$$\frac{n}{N}\{\hat{R}_n(\hat{F}_{\hat{m}}) - \hat{R}_n(F^*)\} < 2e\beta_\Lambda \left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}} + \nu_n^2 e^{\Psi_n}. \tag{3.8}$$

Remark. The last term in (3.8) suggests that the role of step-size restrictions is to prevent the curvature of a generic risk functional from derailing convergence: Recall from (2.18) that e^{Ψ_n} describes the curvature of $\hat{R}_n(\hat{F}_m)$. For a strongly convex risk with bounded curvature, Lemmas 4.2 and A.1 in [58] show that convergence can still be achieved with unrestricted greedy step-sizes. The corresponding expression is however quite complicated and obscures the relationship between step-size and curvature.

Remark. Regardless of whether the risk curvature is bounded or not, smaller step-sizes always improve the convergence bound. This can be seen from the parsimonious relationship between ν_n and (3.8). Fixing n, pushing the value of ν_n down towards zero yields the unattainable lower limit

$$2e\beta_\Lambda \left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}}.$$

This matches the empirical findings in [15], which reports diminishing gains from shrinkage at the expense of more iterations.

Guarantees. For the propositions below we assume that $n \geq 55$ and set $\nu_n^2 e^{\Psi_n} = \log n/(64n^{1/4})$. As a reminder, $\hat{\lambda}_{boost} = e^{\hat{F}_{\hat{m}}}$ and $\lambda^* = e^{F^*}$; $\kappa > 0$ is a universal constant; $\alpha_{\mathcal{F}}$ and β_Λ come from (2.7) and (2.15) respectively; the complexity measure $J_{\mathcal{F}_\epsilon}$ and the gap $\hat{\gamma}$ are defined in (3.4) and (3.7).

Proposition 5. *(Consistency under correct model specification) Suppose $\mathcal{F}$ contains the true log-intensity function $\log\lambda$. Then with probability*

$$1 - 8\exp\left\{-\left(\frac{\log n}{\kappa\alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon} e^{\|\log\lambda\|_\infty}}\right)^2\right\}$$

we have

$$\left\|\hat{F}_{\hat{m}} - \log\lambda\right\|_{\mu,2}^2 < \frac{13\beta_\Lambda}{\min_{t,x}(\lambda \wedge \hat{\lambda}_{boost})}\left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}}.$$

Remark. In view of the norm equivalence relations (2.8), the result implies that $\hat{F}_{\hat{m}}$ also converges uniformly to $\log\lambda$ in probability.

Remark. As alluded to in the discussion preceding Lemma 3, the performance tradeoff from using low complexity ϵ-gradients is captured above: ϵ and $J_{\mathcal{F}_\epsilon}$ tend to have an inverse relationship, but the bound is minimized by $\epsilon \to 0$ while the probability of it holding is maximized by $J_{\mathcal{F}_\epsilon} \to 0$. Beyond this qualitative structure, our framework is too crude to permit analytic determination of the optimal tradeoff. Instead, cross-validation will be used in the applications in section 4.

Proposition 6. *(Oracle inequality for regression trees) Suppose $e^F \in \mathcal{F}$ for all $F \in \mathcal{F}$. Then with probability*

$$1 - 8\exp\left\{-\left(\frac{\log n}{\kappa\alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon} e^{\|F^*\|_\infty}}\right)^2\right\}$$

we have

$$\left\|\hat{\lambda}_{boost} - \lambda\right\|_{\mu,2}^2 < \min_{F\in\mathcal{F}}\left\|e^F - \lambda\right\|_{\mu,2}^2 + 13\beta_\Lambda\frac{\max_{t,x}(\lambda^* \vee \hat{\lambda}_{boost})^2}{\min_{t,x}(\lambda^* \wedge \hat{\lambda}_{boost})}\left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}}.$$

4 Applications

We apply Algorithm 2 to two simulated datasets as well as an actual one on patient service times at an academic emergency department in the United States. We take $\mathcal{F}$ to be the span of the collection of histogram functions $\phi_j(t,x) = I_{B_j}(t,x)$ where $B_1, \cdots, B_d$ are disjoint cubes indexed by $j = (j_0, j_1, \cdots, j_p)$:

$$B_j = \left\{ (t,x) \quad : \quad \begin{array}{c} t^{(j_0)} < t \le t^{(j_0+1)} \\ x^{(1,j_1)} < x^{(1)} \le x^{(1,j_1+1)} \\ \vdots \\ x^{(p,j_p)} < x^{(p)} \le x^{(p,j_p+1)} \end{array} \right\}, \tag{4.1}$$

with the requirement that $[0,1]\times\mathcal{X} \subseteq \cup_{j=1}^d B_j$. The space $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$ then consists of flexible piecewise constant functions of the form $F(t,x) = \sum_{j=1}^d c_{F,j} I_{B_j}(t,x)$

with orthonormal basis $\left\{I_{B_j}(t,x)/\hat{\mu}_n(B_j)^{1/2} : \hat{\mu}_n(B_j) > 0\right\}$. Note that $\mathcal{F}$ satisfies the exponential closure condition in Proposition 6, and the expressions for $\hat{R}_n(F)$ and $\hat{g}_F$ in Lemma 1 simplify to

$$\hat{R}_n(F) = \sum_{j=1}^{d} \left(e^{c_{F,j}} \hat{\mu}_n(B_j) - \frac{c_{F,j}\widehat{FAIL}_j}{N} \right), \tag{4.2}$$

$$\hat{g}_F(t,x) = \begin{cases} e^{c_{F,j}} - \frac{\widehat{FAIL}_j}{N\hat{\mu}_n(B_j)} & (t,x) \in B_j \text{ and } \hat{\mu}_n(B_j) > 0 \\ \text{free} & (t,x) \in B_j \text{ and } \hat{\mu}_n(B_j) = 0 \end{cases}, \tag{4.3}$$

where $\widehat{FAIL}_j$ is the number of observed failures that occurred in the time-covariate region B_j. That is, $\widehat{FAIL}_j = \sum_i \Delta_i I[\{T_i, X_i(T_i)\} \in B_j]$. Since $\hat{g}_F$ is undefined in certain regions, some smoothing is needed: The descent direction $-\hat{g}_F^{\epsilon}$ is obtained from fitting a limited depth regression tree (Breiman et al. [8]) to $-\hat{g}_F$ using the weighted least squares criterion

$$\sum_{j:\hat{\mu}_n(B_j)>0} \hat{\mu}_n(B_j) \left\{ \hat{g}_F(t_{B_j}, x_{B_j}) - \hat{g}_F^{\epsilon}(t_{B_j}, x_{B_j}) \right\}^2,$$

where (t_{B_j}, x_{B_j}) denotes the centre of B_j. Note that $\hat{g}_F^{\epsilon} \in \mathcal{F}$ because each of its piecewise constant region is the disjoint union of a subset of B_js. In setting the regularization parameters, the number of tree splits (between two to seven) and the number of descent iterations $\hat{m}$ (among 50, 100, 200, 500, 750, and 1,000) are determined using ten-fold cross validation. While Lemma 4 shows that a smaller shrinkage factor ν_n is always better, this also increases the number of iterations. For the examples considered here, setting $\nu_n = 1$ appears to suffice for sample sizes up to $100,000$. Rather than use the simplified step-size $\nu_n/(m+1)$ in Algorithm 2, we perform line-search within the interval $(0, \nu_n/(m+1)]$ to speed up convergence.

4.1 Simulation of Weibull mixture

We simulate $N = 100,000$ observations from the following univariate intensity model:

$$\lambda(t,x) = \begin{cases} 9t/2 & 1 \le x \le 61 \\ 81t^2/8 & 61 < x \le 100 \end{cases}. \tag{4.4}$$

The surface of λ is given in the top panel of Figure 4.1. Each covariate trajectory $X_i(t)$ is piecewise constant, and $X_i(0)$ is drawn uniformly from the integers 1 to 100. After every 0.01 time units, $X_i(t)$ is randomly incremented by one of $\{-1, 0, +1\}$. At the boundaries $x = 1$ or $x = 100$ the increment sets are $\{0, +1\}$ or $\{-1, 0\}$ respectively. Observations that have not yet failed by time 1 are censored (approximately 8%).

For the boundaries that define (4.1) we partition the time interval $[0, 1]$ and $\mathcal{X} = (0, 100]$ each into 100 equal divisions, *i.e.* $d = 10,000$. The bottom left panel of Figure 4.1 shows the intensity estimator $\hat{\lambda}_{boost} = e^{\hat{F}_{\hat{m}}}$. The bottom right panel displays the exact MLE, which shows clear signs of overfitting. This illustrates the need for regularization like those employed in boosting.

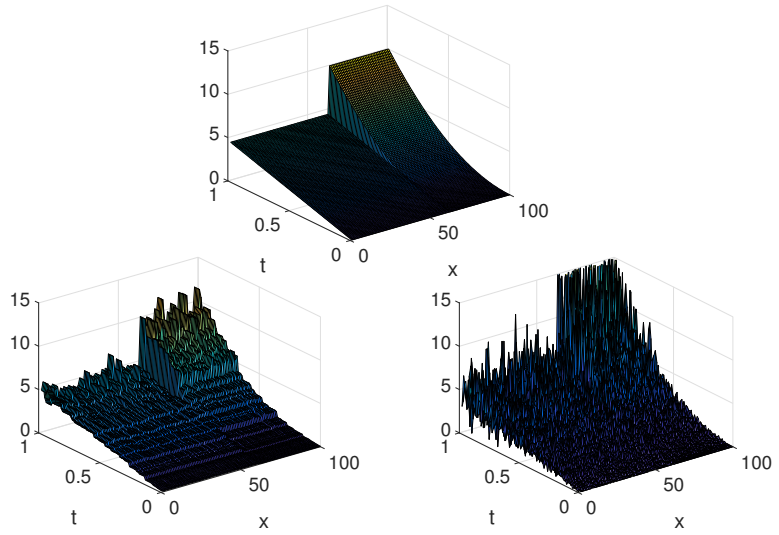

Figure 4.1: Estimation results for the intensity (4.4). *Top*: True intensity function. *Bottom left*: Intensity estimator $\hat{\lambda}_{boost} = e^{\hat{F}_{\hat{m}}}$; two splits used for tree fits, number of iterations $\hat{m} = 200$. *Bottom right*: The exact MLE for exponentially closed classes is given by $\hat{\lambda}$ in Lemma 1.

4.2 Simulation of non-IID covariates

In section 3 we derived guarantees for our estimator under the IID setting[7]. However, this might not hold in every application. Consider the estimation of queuing service rates in the next subsection: Suppose the speed with which a customer is served varies with the number of other customers concurrently served by the same processor. Then customers who overlap may have related processor workload covariates. Rather than subselecting customers who are sufficiently separated in time for the analysis, we assess the robustness of our procedure to potentially non-IID covariates by means of a queuing simulation. The intensity estimated from the simulated data appears to be consistent with the actual intensity used to generate the simulation.

The details of the simulation are as follows. Customer arrivals are Poisson with rate $1/2$ per hour from midnight to noon, and with rate 20 from noon to midnight. Upon arrival a customer is revealed to be of type $x^{(2)} = 1$ or $x^{(2)} = 2$ with equal probability. The queue has a server that can serve up to three customers simultaneously, but if it is fully loaded then new customers are queued on a first-come-first-serve basis. The service rate for a customer at time t depends on the total number of customers currently being served

[7] This setting enables the use of standard concentration results in Lemma 3 and Propositions 5-6.

$x^{(1)}(t) \in \{1, 2, 3\}$ as well as her type $x^{(2)}$ in the following way:

$$\lambda(t, x^{(1)}, x^{(2)}) = \begin{cases} \frac{3}{2} - \frac{1}{2x^{(1)}} - \frac{3}{4}(t - \frac{1}{2}) & x^{(2)} = 1 \\ \frac{1}{2} + \frac{1}{2x^{(1)}} + \frac{1}{2}(t - \frac{1}{2}) & x^{(2)} = 2 \end{cases} . \tag{4.5}$$

Thirty thousand service completions are simulated from this system, and customers whose service duration exceed one hour are censored (approximately one third). For the boundaries that define (4.1) we partition the time interval into 100 equally spaced divisions and note that both covariates are categorical. Figure 4.2 displays the estimation results for each covariate pair combination $(x^{(1)}, x^{(2)})$. The solid line denotes the true intensity, which is tracked by the piecewise constant one representing our estimate, despite $x^{(1)}$ not being IID. For comparison, we also conduct parametric hazard regressions on the data and find that the exponential and the log-logistic distributions provide the best fits. It is clear from the figure that neither are able to track the way the rate changes as the covariates vary.

4.3 Actual patient service times in an emergency department

A series of recent empirical papers in the operations discipline examine the impact of staff workload on service rates in hospitals [2, 3, 4, 5, 6, 13, 25, 26, 27, 31]. A wide spectrum of relationships have been reported ranging from slow-down [3, 36], to no effect [36, 40], to speed-up [2, 25, 26, 27], to slow-down and then speed-up [4, 31], and to still more complicated patterns [6]. These studies utilize a separable covariate specification for estimation. An intriguing question is whether it might be that all these effects actually coexist within the same service process: If the process depends on nontrivial interactions between workload and the other covariates, then a different qualitative impact might be observed for a different combination of the other covariates. The reported findings might then represent different projections of the truth onto various separable specifications. Furthermore, if the datasets used in these studies differ in patient case mix, then further heterogeneity may arise from the projection of the interactions. We test the plausibility of this interaction mechanism by using our procedure to explore covariate interactions in a real dataset for emergency department (ED) service times. Comparing our estimate against a separable specification that is fit to the same dataset yields results that are consistent with this hypothesis.

The data comes from the ED of an academic hospital in the United States and is kindly provided by Dr. Kito Lord. It contains information on 86,983 treatment encounters from 2014 to Q1 of 2015. For each encounter the dataset records age and gender, the times of ward entry and departure (treatment commencement and discharge), and the assigned bedspace number. In addition, the Emergency Severity Index (ESI) [1] of each patient is also recorded, with level 1 being the most severe (e.g., cardiac arrest) and level 5 the least (e.g., rash), although the former group of trauma cases is excluded from this analysis.

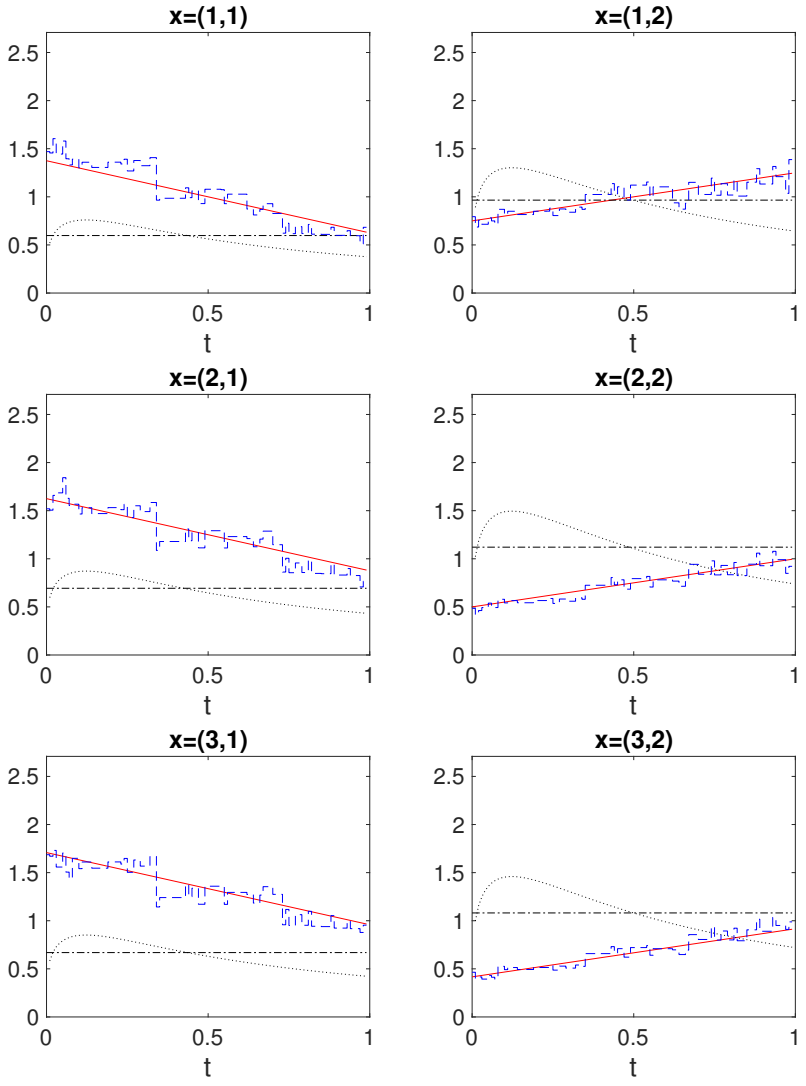

Figure 4.2: Estimation results for the ED simulation (4.5). The estimated intensity for each covariate pair combination $(x^{(1)}, x^{(2)})$ is illustrated in a separate panel. The red solid line denotes the true intensity, which is correctly tracked by the piecewise constant one representing $\hat{\lambda}_{boost}$. Two splits are used for the tree fits, $\hat{m} = 50$. The horizontal dash-dot line represents the parametric exponential fit, and the dotted curve the log-logistic fit.

Time	ESI	Age	Census	Time of day	Nurseload	Day of week	Gender
1	0.7	0.03	0.003	0.001	0.0009	0.0004	0.0002

Table 1: The relative importances of time dependence and the covariates for the real ED dataset. The numbers are scaled so that the largest value is 1.

We extract from the dataset the following covariates for use in estimating the treatment service rate in the ED ward: Time-fixed covariates include age, gender, ESI, time of day at ward entry, and day of week. Time-varying ones include the ward census (total number of occupied beds) and nurseload (number of patients seen by the nurse treating the individual[8]) time series over the course of treatment. These serve as measures of staff workload. A recent paper [5] also examines the impact of shift handovers together with load, but our dataset does not contain shift information so we are unable to incorporate this factor. If shift data is available then we can also study whether subjecting staff to long periods of high load slows down service [16, 25].

For the boundaries that define (4.1) we partition the time interval [0,10] hours into 30 minute buckets. Service times beyond 10 hours are censored (less than 2% of encounters). The boundaries for the age dimension are given by the partitions $\leq$30 years, (30,39], (39,49], (49,59], (59,69], (69,79], and >79. The boundaries for time of day at entry are set two hours apart, and ward census is split according to $\leq$14 occupied beds, (14,29], (30,44], (45,59], and >59. Cross validation results point to using three tree splits and $\hat{m} = 500$ iterations, which suggests the presence of third order interactions among time and the covariates. The relative importances [15] of the variables in the boosted tree estimator are given in Table 1. By far the most influential factor affecting the estimated service rate is time since treatment commencement, followed by ESI level. In Figure 4.3 we profile the rate by time, ESI level, and ward census while holding the other covariates fixed at their median values.

Two features stand out from Figure 4.3. First, by fixing a census level we see that the service rate is unimodal in time, with shapes that resemble hazard functions belonging to a family of log-normal or log-logistic distributions. This agrees with the literature that find log-normality to be a reasonable parametric fit for service durations [3, 9, 20]. Second, by fixing a timepoint we see that service first speeds up and then slows down as census grows. However this relationship is not the only one that exists in our data, as evidenced by the interaction of census with the other factors: The profile plots displayed in Figure 4.4 are for ESI level 3 with the other covariates fixed at their lower quartile values (left panel) and at their upper quartile values (right panel). For the lower quartile configuration we recover the finding in [3] that service actually slows down as census increases, and this relationship holds for other ESI levels as well.

Given the presence of more than one type of workload effect in the data when

[8] Computed from the fact that each bedspace subgroup is assigned to one nurse, so it suffices to know which beds are occupied at any point in time.

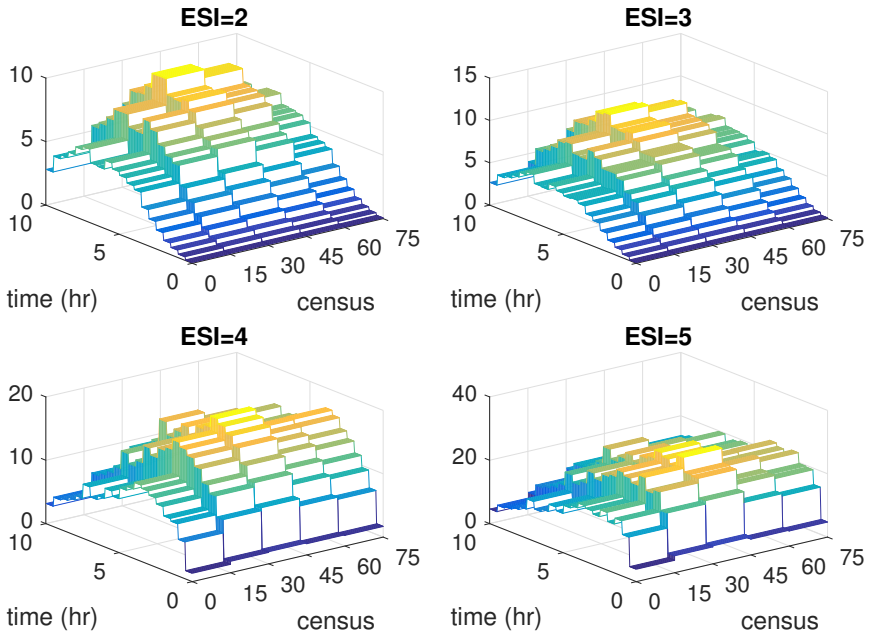

Figure 4.3: Profile of $\hat{\lambda}_{boost}$ by time, ESI, and dynamic ward census. All other covariates are fixed at their median values.

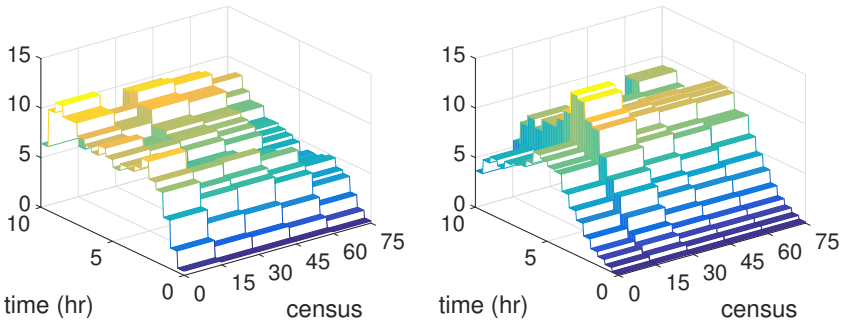

Figure 4.4: Profiles of $\hat{\lambda}_{boost}$ by time and ward census for ESI level 3. *Left*: All other covariates are fixed at their lower quartile values. *Right*: All other covariates are fixed at their upper quartile values.

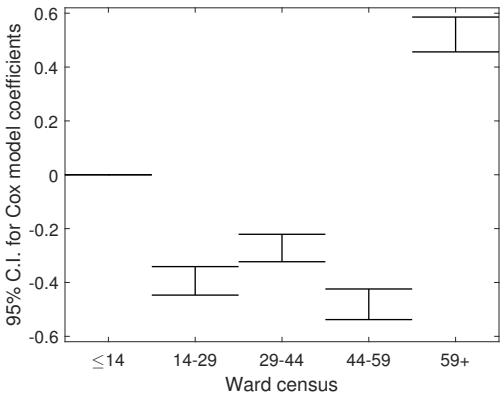

Figure 4.5: Confidence intervals for the Cox model coefficients of the dynamic ward census levels. The lowest level (≤ 14 occupied beds) serves as baseline and is hence set to zero.

interactions are accounted for, it is natural to ask what the estimated effect looks like under a separable covariate specification like those used in literature. We first transform each continuous covariate into a categorical one using the boundaries mentioned earlier as breakpoints. A Cox proportional hazards model for time-varying covariates is then fit to these non-interacted dummy variables. Figure 4.5 graphs the confidence intervals for the coefficients of different census levels. We see that the separable specification matches the findings in [4, 31]: Service initially slows down and then speeds up as census increases, which is once again different from what we have seen so far. Taken together, our results suggest that first order non-interacted effect specifications might only paint a partial picture of the relationship between workload and service: Higher order interactions appear to be important in this setting, and our procedure now provides a way for operations scholars to investigate these interactions in more depth.

5 Discussion

For a given application, the proposed estimator can potentially be used to evaluate the goodness-of-fit of simpler parametric hazard models. Since our approach is likelihood-based, future work might examine whether model selection frameworks like those in Vuong [57] can be extended to cover partial likelihood functionals. For this, Bühlmann and Hothorn [10] and the discussion by Hastie [19] provide some guidance for determining the effective degrees of freedom for the boosting estimator. The ideas in Zou et al. [59], Tibshirani and Taylor [54], and Tibshirani [53] may also be germane.

While the guarantees for our estimator are derived under the IID setting, simulation results tentatively suggest that our procedure might also work for

non-IID covariates that arise from queuing contexts. We conjecture that it is possible to relax the IID assumption to weakly-dependent covariate sequences by following Lozano et al. [35], in which such an analysis is demonstrated for classification losses.

Acknowledgements

We thank Trevor Hastie for suggesting the use of a likelihood functional approach for this problem and for his valuable feedback. In addition, we are grateful to Brian Clarke, Jack Hall, Sahand Negahban, and Anru Zhang for helpful discussions. The dataset analyzed in section 4 is kindly provided by Dr. Kito Lord.

Appendix: Proofs

Proof of Lemma 1. Fix a realization of $\{L_i, U_i, C_i, X_i(\cdot)\}_{i=1}^N$. Using (2.4) we can rewrite (1.1) as

$$\hat{R}_n(F) = \int e^F d\hat{\mu}_n - \frac{1}{n}\sum_{i=1}^n \Delta_i F(T_i, X_i(T_i)).$$

The last term equals $-\int \hat{\lambda} F d\hat{\mu}_n$ by the orthonormality of $\{\hat{\varphi}_j\}_j$, which completes the derivation of (2.9).

By an interchange argument we obtain

$$\begin{aligned} \frac{d}{d\theta}\hat{R}_n(F+\theta f) &= \int (e^{F+\theta f} - \hat{\lambda}) f d\hat{\mu}_n, \\ \frac{d^2}{d\theta^2}\hat{R}_n(F+\theta f) &= \int e^{F+\theta f} f^2 d\hat{\mu}_n, \end{aligned} \tag{5.1}$$

the latter being positive whenever $f \neq 0$, *i.e.* $\hat{R}_n(F)$ is convex. The Taylor representation (2.10) then follows from noting that $\hat{g}_F$ is the orthogonal projection of $e^F - \hat{\lambda} \in L^2(\hat{\mu}_n)$ onto $(\mathcal{F}, \langle\cdot,\cdot\rangle_{\hat{\mu}_n})$. □

Proof of Lemma 2. Writing

$$R(F) = \frac{1}{N}\sum_{i=1}^N \mathbb{E}\left\{\int_0^1 I(L_i \leq t \leq T_i) e^{F(t, X_i(t))} dt - \Delta_i F(T_i, X_i(T_i))\right\},$$

we can apply (2.3) to establish the first part of the integral in (2.12) when $F \in \mathcal{F} \cup \{\log\lambda\}$. To complete the representation it suffices to show that the point process $M(B) = \Delta \cdot I[\{T, X(T)\} \in B]$ has mean $\int_B \lambda d\mu$. To this end recall from the introduction that the covariate process is assumed to be external to the failure time U, and that U and (L, C) are conditionally independent

given $X(\cdot)$. Hence the density of U conditional on $\{L, C, X(\cdot)\}$ is $\mathbb{P}\{U \geq t \mid X(\cdot)\}\lambda(t-, X(t-))$. Under A3, the expectation of $M(B)$ exists and

$$\begin{aligned}
&\mathbb{E}(\Delta \cdot I[\{T, X(T)\} \in B]) \\
=&\mathbb{E}\mathbb{E}\left\{I(L \leq U \leq C) \cdot I[\{U, X(U)\} \in B] \mid L, C, X(\cdot)\right\} \\
=&\mathbb{E}\int_0^1 I(L \leq t \leq C)\mathbb{P}\{U \geq t \mid X(\cdot)\} \cdot \lambda(t, X(t)) I[\{t, X(t)\} \in B] dt \\
=&\mathbb{E}\int_0^1 \mathbb{E}\left\{I(L \leq t \leq C) I(t \leq U) \cdot \lambda(t, X(t)) I[\{t, X(t)\} \in B] \mid L, C, X(\cdot)\right\} dt \\
=&\mathbb{E}\int_0^1 \mathbb{E}\left\{I(L \leq t \leq T) \cdot \lambda(t, X(t)) I[\{t, X(t)\} \in B] \mid L, C, X(\cdot)\right\} dt.
\end{aligned}$$

By a Fubini type argument for conditional expectations we further obtain

$$\begin{aligned}
&\mathbb{E}(\Delta \cdot I[\{T, X(T)\} \in B]) \\
=&\mathbb{E}\mathbb{E}\left\{\int_0^1 I(L \leq t \leq T) \cdot \lambda(t, X(t)) I[\{t, X(t)\} \in B] dt \,\middle|\, L, C, X(\cdot)\right\} \\
=&\mathbb{E}\int_0^1 I(L \leq t \leq T) \cdot \lambda(t, X(t)) I[\{t, X(t)\} \in B] dt \\
=&\int_B \lambda d\mu,
\end{aligned}$$

where the last equality follows from A1 and (2.3).

Now note that for a positive constant Λ the function $e^y - \Lambda y$ is bounded below by both $-\Lambda y$ and $\Lambda y + 2\Lambda\{1 - \log 2\Lambda\}$, hence $e^y - \Lambda y \geq \Lambda|y| + 2\Lambda \min\{0, 1 - \log 2\Lambda\}$. Since $\Lambda \min\{0, 1 - \log 2\Lambda\}$ is non-increasing in Λ, A1 implies that

$$\begin{aligned}
e^{F(t,x)} - \lambda(t,x)F(t,x) &\geq \min\left\{e^{F(t,x)} - \Lambda_L F(t,x), e^{F(t,x)} - \Lambda_H F(t,x)\right\} \\
&\geq \Lambda_L |F(t,x)| + 2\Lambda_H \min\{0, 1 - \log(2\Lambda_H)\}.
\end{aligned}$$

Integrating both sides and using the norm equivalence relation (2.8) shows that

$$\begin{aligned}
R(F) &\geq \Lambda_L \|F\|_{\mu,1} + 2\Lambda_H \min\{0, 1 - \log(2\Lambda_H)\} \\
&\geq \frac{2\Lambda_L}{\alpha_{\mathcal{F}}} \|F\|_\infty + 2\Lambda_H \min\{0, 1 - \log(2\Lambda_H)\} \\
&\geq \frac{2\Lambda_L}{\alpha_{\mathcal{F}}} \|F\|_{\mu,2} + 2\Lambda_H \min\{0, 1 - \log(2\Lambda_H)\}.
\end{aligned}$$

The lower bound (2.13) then follows from the second inequality. The last inequality shows that $R(F)$ is coercive on $(\mathcal{F}, \langle\cdot,\cdot\rangle_\mu)$. Moreover the same argument used to derive (2.10) shows that $R(F)$ is smooth and convex on $(\mathcal{F}, \langle\cdot,\cdot\rangle_\mu)$. Therefore a unique minimizer F^* of $R(F)$ exists in $(\mathcal{F}, \langle\cdot,\cdot\rangle_\mu)$. Since A4 implies there is a bijection between the equivalent classes of $(\mathcal{F}, \langle\cdot,\cdot\rangle_\mu)$ and the functions in $\mathcal{F}$, F^* is also the unique minimizer of $R(F)$ in $\mathcal{F}$. Finally, since $e^{F(t,x)} - \lambda(t,x)F(t,x)$ is pointwise bounded below by $\lambda(t,x)\{1 - \log\lambda(t,x)\}$, $R(F) \geq \int(\lambda - \lambda\log\lambda)d\mu = R(\log\lambda)$ for all $F \in \mathcal{F}$. □

Proof of Lemma 3. By a pointwise-measurable argument (Example 2.3.4 of [55]) it can be shown that all suprema quantities appearing below are sufficiently well behaved, so outer integration is not required. Suppose the Orlicz norm $|||X||| = \inf\{C > 0 : \mathbb{E}e^{(X/C)^2} \leq 2\}$ for the following quantities satisfy

$$\left|\left|\left| \sup_{F \in \mathcal{F}_\epsilon^{\Psi_n}} \left| \frac{n}{N}\{\hat{R}_n(F) - \hat{R}_n(0)\} - \{R(F) - R(0)\} \right| \right|\right| \leq \kappa' J_{\mathcal{F}_\epsilon}/n^{1/4}, \tag{5.2}$$

$$\left|\left|\left| \sup_{G \in \mathcal{F}_\epsilon : \|G\|_\infty = 1} \left| \frac{n}{N}\|G\|_{\hat{\mu}_n,1} - \|G\|_{\mu,1} \right| \right|\right| < \kappa'' J_{\mathcal{F}_\epsilon}/n^{1/2}, \tag{5.3}$$

where $J_{\mathcal{F}_\epsilon}$ is the complexity measure (3.4), and κ', κ'' are universal constants. Then by Markov's inequality, (3.6) holds with probability $1-2\exp[-\{\eta n^{1/4}/(\kappa' J_{\mathcal{F}_\epsilon})\}^2]$, and

$$\sup_{G \in \mathcal{F}_\epsilon : \|G\|_\infty = 1} \left\{ \|G\|_{\mu,1} - \frac{n}{N}\|G\|_{\hat{\mu}_n,1} \right\} < 1/\alpha_{\mathcal{F}} \tag{5.4}$$

holds with probability $1-2\exp[-\{n^{1/2}/(\alpha_{\mathcal{F}}\kappa'' J_{\mathcal{F}_\epsilon})\}^2]$. Since $\alpha_{\mathcal{F}} > 1$ and $\eta < 1$, (3.6) and (5.4) jointly hold with probability $1-4\exp[-\{\eta n^{1/4}/(\kappa \alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon})\}^2]$. The lemma then follows if (5.4) implies (3.5). Indeed, for any non-zero $F \in \mathcal{F}$, its normalization $G = F/\|F\|_\infty$ is in $\mathcal{F}_\epsilon$ by construction (3.3). Then (5.4) implies that

$$\frac{\|F\|_\infty}{\|F\|_{\hat{\mu}_n,1}} = 1/\|G\|_{\hat{\mu}_n,1} \leq \alpha_{\mathcal{F}}$$

because

$$1/\alpha_{\mathcal{F}} > \|G\|_{\mu,1} - \frac{n}{N}\|G\|_{\hat{\mu}_n,1} \geq 2/\alpha_{\mathcal{F}} - \|G\|_{\hat{\mu}_n,1},$$

where the last inequality follows from the definition of $\alpha_{\mathcal{F}}$ (2.7).

Thus it remains to establish (5.2) and (5.3), which is an exercise in applying the symmetrization and maximal inequality results in sections 2.2 and 2.3.2 of [55]. Write $\hat{R}_n(F) = (1/n)\sum_{i=1}^N l_i(F)$ where $l_i(F) = \int_{L_i}^{T_i} e^{F(t,X_i(t))}dt - \Delta_i F(T_i, X_i(T_i))$ are independent copies of the loss

$$l(F) = \int_0^1 I(L \leq t \leq T) e^{F(t,X(t))}dt - \Delta \cdot F(T, X(T)), \tag{5.5}$$

which is a stochastic process indexed by $F \in \mathcal{F}$. As was shown in Lemma 2, $\mathbb{E}l(F) = R(F)$. Let $\varsigma_1, \cdots, \varsigma_N$ be independent Rademacher random variables that are independent of $Z = \{L_i, U_i, C_i, X_i(\cdot)\}_{i=1}^N$. It follows from the symmetrization Lemma 2.3.6 of [55] for stochastic processes that the left hand side

of (5.2) is bounded by twice the Orlicz norm of

$$\begin{aligned}
&\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}\left|\frac{1}{N}\sum_{i=1}^N \varsigma_i\{l_i(F)-l_i(0)\}\right| \\
\leq&\frac{1}{N}\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}\left|\sum_{i=1}^N \varsigma_i\int_{L_i}^{T_i}\left\{e^{F(t,X_i(t))}-1\right\}dt\right| \\
+&\frac{1}{N}\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}\left|\sum_{i=1}^N \varsigma_i\Delta_i F(T_i,X_i(T_i))\right|.
\end{aligned} \tag{5.6}$$

Now hold Z fixed so that only $\varsigma_1,\cdots,\varsigma_N$ are stochastic, in which case the sum in the second line of (5.6) becomes a separable subgaussian process. Since the Orlicz norm of $\sum_{i=1}^N \varsigma_i a_i$ is bounded by $(6\sum_{i=1}^N a_i^2)^{1/2}$ for constants a_i, the sum satisfies the Lipschitz property

$$\begin{aligned}
&\left\|\sum_{i=1}^N \varsigma_i\int_{L_i}^{T_i}\left\{e^{F(t,X_i(t))}-e^{F'(t,X_i(t))}\right\}dt\right\|_{\varsigma|Z}^2 \\
\leq&6\sum_{i=1}^N\left[\int_{L_i}^{T_i}\left\{e^{F(t,X_i(t))}-e^{F'(t,X_i(t))}\right\}dt\right]^2 \\
\leq&6e^{2\Psi_n}\sum_{i=1}^N\left(\int_{L_i}^{T_i}|F(t,X_i(t))-F'(t,X_i(t))|dt\right)^2 \\
\leq&6e^{2\Psi_n}\sum_{i=1}^N\int_{L_i}^{T_i}\{F(t,X_i(t))-F'(t,X_i(t))\}^2dt \\
=&6ne^{2\Psi_n}\|F-F'\|_{\hat{\mu}_n,2}^2,
\end{aligned}$$

where the second inequality follows from $|e^x-e^y|\leq e^{\max(x,y)}|x-y|$ and the last from the Cauchy-Schwarz inequality. Putting the Lipschitz constant $(6n)^{1/2}e^{\Psi_n}$ obtained above into Theorem 2.2.4 of [55] yields the following maximal inequality: There is a universal constant κ' such that

$$\begin{aligned}
&\left\|\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}\left|\sum_{i=1}^N \varsigma_i\int_{L_i}^{T_i}\left\{e^{F(t,X_i(t))}-1\right\}dt\right|\right\|_{\varsigma|Z} \\
\leq&\kappa' n^{1/2}e^{\Psi_n}\int_0^{\Psi_n}\left\{\log\mathcal{N}(u,\mathcal{F}_\epsilon^{\Psi_n},\hat{\mu}_n)\right\}^{1/2}du \\
\leq&\kappa' n^{1/2}e^{\Psi_n}\Psi_n J_{\mathcal{F}_\epsilon},
\end{aligned}$$

the last line following from (3.4). Likewise the conditional Orlicz norm for the supremum of $\left|\sum_{i=1}^N \varsigma_i\Delta_i F(T_i,X_i(T_i))\right|$ is bounded by $\kappa' J_{\mathcal{F}_\epsilon} n^{1/2}\Psi_n$. Since

neither bounds depend on Z, plugging back into (5.6) proves (5.2):

$$\begin{aligned}&\left\|\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}\left|\frac{n}{N}\{\hat{R}_n(F)-\hat{R}_n(0)\}-\{R(F)-R(0)\}\right|\right\|\\ \leq&2\kappa' J_{\mathcal{F}_\epsilon}\frac{\Psi_n e^{\Psi_n}}{n^{1/2}}\left\{1+e^{-\Psi_n}\right\}\leq 4\kappa' J_{\mathcal{F}_\epsilon}/n^{1/4},\end{aligned}$$

where $\Psi_n e^{\Psi_n}=n^{1/4}$ by definition (2.17). On noting that $\|G\|_{\hat{\mu}_n,1}$ equals $(1/n)\sum_{i=1}^N\int_{L_i}^{T_i}|G(t,X_i(t))|dt$ and $\mathbb{E}\int_L^T|G(t,X(t))|dt=\|G\|_{\mu,1}$, (5.3) can be established using the same approach. □

Proof of Lemma 4. For $m<\hat{m}$, applying (2.10) and (2.16) to $\hat{R}_n(\hat{F}_{m+1})=\hat{R}_n(\hat{F}_m-\frac{\nu_n}{m+1}\hat{g}^\epsilon_{\hat{F}_m})$ yields

$$\begin{aligned}\hat{R}_n(\hat{F}_{m+1})=&\hat{R}_n(\hat{F}_m)-\frac{\nu_n}{m+1}\left\langle\hat{g}_{\hat{F}_m},\hat{g}^\epsilon_{\hat{F}_m}\right\rangle_{\hat{\mu}_n}\\ &+\frac{\nu_n^2}{2(m+1)^2}\int(\hat{g}^\epsilon_{\hat{F}_m})^2\exp\left\{\hat{F}_m+\hat{\rho}_m(\hat{F}_{m+1}-\hat{F}_m)\right\}d\hat{\mu}_n\\ <&\hat{R}_n(\hat{F}_m)-\frac{(1-\epsilon)\nu_n}{m+1}\|\hat{g}_{\hat{F}_m}\|_{\hat{\mu}_n,2}+\frac{\nu_n^2e^{\Psi_n}}{2(m+1)^2},\end{aligned}\tag{5.7}$$

where the bound for the integral follows from the fact that $\max_{m\leq\hat{m}}\|\hat{F}_m\|_\infty<\Psi_n$ under (2.19), and $\int(\hat{g}^\epsilon_{\hat{F}_m})^2d\hat{\mu}_n=1$ by definition (2.16). Then in particular for $m\leq\hat{m}$,

$$\hat{R}_n(\hat{F}_m)<\hat{R}_n(0)+\sum_{m=0}^\infty\frac{\nu_n^2e^{\Psi_n}}{2(m+1)^2}<\hat{R}_n(0)+1\leq 2.$$

Since $\max_{m\leq\hat{m}}\|\hat{F}_m\|_\infty<\Psi_n$, by the assumption $\sup_{F\in\mathcal{F}_\epsilon^{\Psi_n}}|(n/N)\hat{R}_n(F)-R(F)|<1$ we have

$$R(\hat{F}_m)\leq\frac{n}{N}\hat{R}_n(\hat{F}_m)+\left|\frac{n}{N}\hat{R}_n(\hat{F}_m)-R(\hat{F}_m)\right|<3.$$

Clearly the minimizer F^* also satisfies $R(F^*)\leq R(0)<3$. Thus coercivity (2.14) implies that

$$\|\hat{F}_m\|_\infty,\|F^*\|_\infty<\alpha_{\mathcal{F}}\beta_\Lambda,\tag{5.8}$$

so the gap $\hat{\gamma}$ defined in (3.7) is bounded as claimed.

It remains to establish (3.8), for which we need only consider the case $\hat{R}_n(\hat{F}_{\hat{m}})-\hat{R}_n(F^*)>0$. The termination criterion $\hat{g}_{\hat{F}_m}=0$ in Algorithm 2 is never triggered under this scenario, for otherwise $\hat{F}_{\hat{m}}$ minimizes $\hat{R}_n(F)$. Thus either $\hat{m}=\infty$ or (2.19) is violated for $m=\hat{m}$:

$$\begin{aligned}\Psi_n\leq&\left\|\hat{F}_{\hat{m}}-\frac{\nu_n}{\hat{m}+1}\hat{g}^\epsilon_{\hat{F}_{\hat{m}}}\right\|_\infty\leq\alpha_{\mathcal{F}}\left\|\hat{F}_{\hat{m}}-\frac{\nu_n}{\hat{m}+1}\hat{g}^\epsilon_{\hat{F}_{\hat{m}}}\right\|_{\hat{\mu}_n,2}\\ \leq&\alpha_{\mathcal{F}}\left(\sum_{m=0}^{\hat{m}-1}\frac{\nu_n}{m+1}+1\right)\end{aligned}\tag{5.9}$$

where the inequalities follow from the hypothesized (3.5), and $\|\hat{g}^{\epsilon}_{\hat{F}_m}\|_{\hat{\mu}_n,2} = 1$. Given the diverging sum in (5.9), the inequality also holds if $\hat{m} = \infty$.

The remark after Lemma 3 implies that F^* lies in the span of $\{\hat{\varphi}_j(t,x)\}_{j=1}^d$, so the representation (2.9) is valid for $\hat{R}_n(F^*)$. Hence by convexity $\left\langle \hat{g}_{\hat{F}_m}, F^* - \hat{F}_m \right\rangle_{\hat{\mu}_n} \leq \hat{R}_n(F^*) - \hat{R}_n(\hat{F}_m)$. Furthermore

$$\begin{aligned}\left\langle \hat{g}_{\hat{F}_m}, \hat{F}_m - F^* \right\rangle_{\hat{\mu}_n} &\leq \|\hat{F}_m - F^*\|_{\hat{\mu}_n,2}\|\hat{g}_{\hat{F}_m}\|_{\hat{\mu}_n,2} \\ &\leq \|\hat{F}_m - F^*\|_{\infty}\|\hat{g}_{\hat{F}_m}\|_{\hat{\mu}_n,2} \leq \hat{\gamma}\|\hat{g}_{\hat{F}_m}\|_{\hat{\mu}_n,2}.\end{aligned}$$

Putting both into (5.7) gives

$$\begin{aligned}\hat{R}_n(\hat{F}_{m+1}) &< \hat{R}_n(\hat{F}_m) + \frac{(1-\epsilon)\nu_n}{m+1}\frac{\left\langle \hat{g}_{\hat{F}_m}, F^* - \hat{F}_m \right\rangle_{\hat{\mu}_n}}{\hat{\gamma}} + \frac{\nu_n^2 e^{\Psi_n}}{2(m+1)^2} \\ &\leq \hat{R}_n(\hat{F}_m) + \frac{1-\epsilon}{\hat{\gamma}} \cdot \frac{\nu_n}{m+1}\{\hat{R}_n(F^*) - \hat{R}_n(\hat{F}_m)\} + \frac{\nu_n^2 e^{\Psi_n}}{2(m+1)^2}.\end{aligned}$$

Subtracting $\hat{R}_n(F^*)$ from both sides above and abbreviating $\hat{\delta}_m = (n/N)\{\hat{R}_n(\hat{F}_m) - \hat{R}_n(F^*)\}$, we obtain

$$\hat{\delta}_{m+1} < \left(1 - \frac{1-\epsilon}{\hat{\gamma}} \cdot \frac{\nu_n}{m+1}\right)\hat{\delta}_m + (n/N)\frac{\nu_n^2 e^{\Psi_n}}{2(m+1)^2}.$$

Since the term inside the first parenthesis is between 0 and 1, solving the recurrence yields

$$\begin{aligned}\hat{\delta}_{\hat{m}} &< \hat{\delta}_0 \prod_{m=0}^{\hat{m}-1}\left(1 - \frac{1-\epsilon}{\hat{\gamma}} \cdot \frac{\nu_n}{m+1}\right) + \nu_n^2 e^{\Psi_n}\sum_{m=0}^{\infty}\frac{1}{2(m+1)^2} \\ &< \max\{0, \hat{\delta}_0\}\exp\left(-\frac{1-\epsilon}{\hat{\gamma}}\sum_{m=0}^{\hat{m}-1}\frac{\nu_n}{m+1}\right) + \nu_n^2 e^{\Psi_n} \\ &< e\max\{0, \hat{\delta}_0\}\exp\left(-\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}\Psi_n\right) + \nu_n^2 e^{\Psi_n},\end{aligned}$$

where in the second inequality we used the fact that $0 \leq 1 + y \leq e^y$ for $|y| < 1$, and the last follows from (5.9).

The Lambert function (2.17) in $\Psi_n = W(n^{1/4})$ is asymptotically $\log y - \log\log y$, and in fact by Theorem 2.1 of Hoorfar and Hassani [21], $W(y) \geq \log y - \log\log y$ for $y \geq e$. Since by assumption $n \geq 55 > e^4$, the above becomes

$$\hat{\delta}_{\hat{m}} < e\max\{0, \hat{\delta}_0\}\left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}} + \nu_n^2 e^{\Psi_n}.$$

The last step is to control $\hat{\delta}_0$, which is bounded by $(n/N)\{1 - \hat{R}_n(F^*)\}$ because $\hat{R}_n(\hat{F}_0) = \hat{R}_n(0) \leq 1$. Then under the hypothesis $|(n/N)\hat{R}_n(F^*) - R(F^*)| < 1$,

we have $\hat{\delta}_0 \leq (n/N) - R(F^*) + 1 < 2 - R(F^*)$. Since (2.13) implies $R(F^*) \geq 2\Lambda_H \min\{0, 1 - \log(2\Lambda_H)\}$,

$$\hat{\delta}_0 < 2 - R(F^*) \leq 2 + 2\Lambda_H \max\{0, \log(2\Lambda_H) - 1\} < 2\beta_\Lambda.$$

□

Proof of Proposition 5. Let $\varepsilon = \log n/(4n^{1/4})$ which is less than one for $n \geq e^4$. Since $\alpha_{\mathcal{F}}, \hat{\gamma} \geq 1$ it follows that

$$\varepsilon < \left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}}. \tag{5.10}$$

Now define the following probability sets

$$S_1 = \left\{ \sup_{F \in \mathcal{F}_\epsilon^{\Psi_n}} \left| \frac{n}{N}\{\hat{R}_n(F) - \hat{R}_n(0)\} - \{R(F) - R(0)\} \right| < \varepsilon/3 \right\},$$

$$S_2 = \left\{ \left| \frac{n}{N}\hat{R}_n(0) - R(0) \right| < \varepsilon/3 \right\},$$

$$S_3 = \left\{ \left| \frac{n}{N}\hat{R}_n(F^*) - R(F^*) \right| < \varepsilon/3 \right\},$$

$$S_4 = \{(3.5) \text{ holds}\},$$

and fix a sample realization from $\cap_{k=1}^4 S_k$. Then the conditions required in Lemma 4 are satisfied with $\sup_{F \in \mathcal{F}_\epsilon^{\Psi_n}} |(n/N)\hat{R}_n(F) - R(F)| < 2\varepsilon/3$, so $\hat{\gamma}$ is bounded and (3.8) holds. Since (2.19) ensures that $\|\hat{F}_{\hat{m}}\|_\infty < \Psi_n$, it also follows that $|(n/N)\hat{R}_n(\hat{F}_{\hat{m}}) - R(\hat{F}_{\hat{m}})| < 2\varepsilon/3$. Combining (3.1) and (3.2) gives

$$\begin{aligned}
\left\|\hat{F}_{\hat{m}} - F^*\right\|_{\mu,2}^2 &\leq \frac{2}{\min_{t,x}(\lambda^* \wedge \hat{\lambda}_{boost})} \left[\varepsilon + \frac{n}{N}\{\hat{R}_n(\hat{F}_{\hat{m}}) - \hat{R}_n(F^*)\}\right] \\
&< \frac{2}{\min_{t,x}(\lambda^* \wedge \hat{\lambda}_{boost})} \left(\frac{\log n}{4n^{1/4}} + 2e\beta_\Lambda \left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}} + \frac{1}{16} \cdot \frac{\log n}{4n^{1/4}}\right) \\
&< \frac{13\beta_\Lambda}{\min_{t,x}(\lambda^* \wedge \hat{\lambda}_{boost})} \left(\frac{\log n}{4n^{1/4}}\right)^{\frac{1-\epsilon}{\alpha_{\mathcal{F}}\hat{\gamma}}},
\end{aligned}$$

where the second inequality follows from (3.8) and $\nu_n^2 e^{\Psi_n} = \log n/(64n^{1/4})$, and the last from (5.10). The desired bound for $\|\hat{F}_{\hat{m}} - \log\lambda\|_{\mu,2}$ then follows from noting from Lemma 2 that $F^* = \log\lambda$.

The next task is to lower bound $\mathbb{P}(\cap_{k=1}^4 S_k)$. It follows from Lemma 3 that

$$\mathbb{P}(S_1 \cap S_4) \geq 1 - 4\exp\left\{-\left(\frac{\log n}{12\kappa\alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon}}\right)^2\right\}.$$

Bounds on $\mathbb{P}(S_2)$ and $P(S_3)$ can be obtained using Hoeffding's inequality. Note from (1.1) that $(n/N)\hat{R}_n(0) = \sum_{i=1}^N \max(0, T_i - L_i)/N$ and $(n/N)\hat{R}_n(F^*) =$

$\sum_{i=1}^{N} l_i(F^*)/N$ for the loss $l(\cdot)$ defined in (5.5). Since $0 \le \max(0, T_i - L_i) \le 1$ and $-\|F^*\|_\infty < l(F^*) \le \|e^{F^*}\|_\infty + \|F^*\|_\infty$,

$$\mathbb{P}(S_2) \ge 1-2\exp\left\{-2n^{1/2}\left(\frac{\log n}{12}\right)^2\right\},\ \mathbb{P}(S_3) \ge 1-2\exp\left\{-2n^{1/2}\left(\frac{\log n}{36e^{\|F^*\|_\infty}}\right)^2\right\}.$$

By increasing the value of κ and/or replacing $J_{\mathcal{F}_\epsilon}$ with $\max(1, J_{\mathcal{F}_\epsilon})$ if necessary, we can combine the inequalities to get a crude but compact bound:

$$\mathbb{P}\{\cap_{k=1}^{4} S_k\} \ge 1 - 8\exp\left\{-\left(\frac{\log n}{\kappa\alpha_{\mathcal{F}} J_{\mathcal{F}_\epsilon} e^{\|F^*\|_\infty}}\right)^2\right\}.$$

□

Proof of Proposition 6. Since F^* minimizes $R(F)$, it satisfies the first order condition $\left\langle e^{F^*} - \lambda, F\right\rangle_\mu = 0$ for all $F \in \mathcal{F}$. Under exponentiation closure $\lambda^* = e^{F^*} \in \mathcal{F}$, so it is the orthogonal projection of λ onto $(\mathcal{F}, \langle\cdot,\cdot\rangle_\mu)$. Hence

$$\begin{aligned}\left\|\hat{\lambda}_{boost} - \lambda\right\|_{\mu,2}^2 &= \left\|e^{F^*} - \lambda\right\|_{\mu,2}^2 + \left\|e^{\hat{F}_{\hat{m}}} - e^{F^*}\right\|_{\mu,2}^2 \\ &\le \min_{F\in\mathcal{F}} \left\|e^F - \lambda\right\|_{\mu,2}^2 + \max_{t,x}(\lambda^* \vee \hat{\lambda}_{boost})^2 \left\|\hat{F}_{\hat{m}} - F^*\right\|_{\mu,2}^2,\end{aligned}$$

where the inequality follows from $|e^x - e^y| \le \max(e^x, e^y)|x - y|$. Bounding the last term in the same way as Proposition 5 yields the desired result. □

References

[1] Emergency severity index (ESI): A triage tool for emergency department. `www.ahrq.gov/professionals/systems/hospital/esi/esi1.html`. Accessed Aug 25 2016.

[2] D. Anderson, C. Price, B. Golden, W. Jank, and E. Wasil. Examining the discharge practices of surgeons at a large medical center. *Health Care Management Science*, 14(4):338–347, 2011.

[3] M. Armony, S. Israelit, Marmor Y. N. Mandelbaum, Y. Tseytlin, and G. B. Yom-Tov. On patient flow in hospitals: A data-based queueing-science perspective. *Stochastic Systems*, 5(1):146–194, 2015.

[4] R. J. Batt and C. Terwiesch. Doctors under load: An empirical study of state-dependent service times in emergency care. 2014.

[5] R. J. Batt, D. S. Kc, B. R. Staats, and B. W. Patterson. The effects of discrete workshifts on non-terminating queues.

[6] J. A. Berry Jaeker and A. L. Tucker. Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Management Science*, Forthcoming.

[7] L. Breiman. Arcing the edge. *U.C. Berkeley Dept. of Statistics Technical Report*, 486, 1997.

[8] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and regression trees*. Chapman and Hall, 1984.

[9] L. Brown, N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, and L. Zhao. Statistical analysis of a telephone call center: A queueing-science perspective. *Journal of the American Statistical Association*, 100(469):36–50, 2005.

[10] P. Bühlmann and T. Hothorn. Boosting algorithms: Regularization, prediction and model fitting. *Statistical Science*, pages 477–505, 2007.

[11] P. Bühlmann and B. Yu. Boosting with the L_2 loss: Regression and classification. *Journal of the American Statistical Association*, 98(462):324–339, 2003.

[12] M. Cule, R. Samworth, and M. Stewart. Maximum likelihood estimation of a multidimensional log-concave density. *Journal of the Royal Statistical Society B*, 72(5):545–607, 2010.

[13] M. Delasay, A. Ingolfsson, B. Kolfal, and K. Schultz. Load effect on service times. 2015.

[14] S. Dohler and L. Ruschendorf. Adaptive estimation of hazard functions. *Probability and Mathematical Statistics*, 22:355–379, 2002.

[15] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29:1189–1232, 2001.

[16] N. Gans, N. Liu, A. Mandelbaum, H. Shen, and H. Ye. Service times in call centers: Agent heterogeneity and learning with some operational consequences. *Institute of Mathematical Statistics*, 6:99–123, 2010.

[17] U. Grenander. On the theory of mortality measurement. *Scandinavian Actuarial Journal*, 2:125–153, 1956.

[18] C. Gu. Penalized likelihood hazard estimation: A general procedure. *Statistica Sinica*, 6:861–876, 1996.

[19] T.J. Hastie. Comment on paper by bühlmann & hothorn. *Statistical Science*, 22:513–515, 2007.

[20] B. He, F. Dexter, A. Macario, and S. A. Zenios. The timing of staffing decisions in hospital operating rooms: Incorporating workload heterogeneity into the newsvendor problem. *M&SOM*, 14(1):99–114, 2012.

[21] A. Hoorfar and M. Hassani. Inequalities on the Lambert W function and hyperpower function. *Journal of inequalities in pure and applied mathematics*, 9(2), 2008.

[22] T. Hothorn, P. Bühlmann, S. Dudoit, A. Molinaro, and M.J. Van Der Laan. Survival ensembles. *Biostatistics*, 7(3):355–373, 2006.

[23] J. Z. Huang and C. J. Stone. The L_2 rate of convergence for event history regression with time-dependent covariates. *Scandinavian Journal of Statistics*, 25:603–620, 1998.

[24] A. F. Karr. Maximum likelihood estimation in the multiplicative intensity model via sieves. *Annals of Statistics*, 15:473–490, 1987.

[25] D. S. Kc and C. Terwiesch. Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Science*, 55(9):1486–1498, 2009.

[26] D. S. Kc and C. Terwiesch. An econometric analysis of patient flows in the cardiac intensive care unit. *M&SOM*, 14(1):50–65, 2012.

[27] S. H. Kim, E. J. Pinker, J. Rimar, and E. H. Bradley. Impact of severity-adjusted workload on health status of patients discharge from an ICU. 2015.

[28] C. Kooperberg and D. B. Clarkson. Hazard regression with interval-censored data. *Biometrics*, 53:1485–1494, 1997.

[29] C. Kooperberg, C. J. Stone, and Y. K. Truong. Hazard regression. *Journal of the American Statistical Association*, 90:78–94, 1995.

[30] C. Kooperberg, C. J. Stone, and Y. K. Truong. The L_2 rate of convergence for hazard regression. *Scandinavian Journal of Statistics*, 2:143–157, 1995.

[31] L. Kuntz, R. Mennicken, and S. Scholtes. *Stress on the ward-An empirical study of the nonlinear relationship between organizational workload and service quality.* Number 277. Ruhr Economic Papers, 2011.

[32] G. Li and H. Doss. An approach to nonparametric regression for life history data using local linear fitting. *Annals of Statistics*, 23(3):787–823, 1995.

[33] H. Li and Y. Luan. Boosting proportional hazards models using smoothing splines, with applications to high-dimensional microarray data. *Bioinformatics*, 21(10):2403–2409, 2005.

[34] O. B. Linton, J. P. Nielsen, and S. van de Geer. Estimating multiplicative and additive hazard functions by kernel methods. *Annals of Statistics*, 31 (2):464–492, 2003.

[35] A. C. Lozano, S. R. Kulkarni, and R. E. Schapire. Convergence and consistency of regularized boosting with weakly dependent observations. *IEEE Trans. Inf. Theory*, 60(1):651–660, 2014.

[36] R. Lucas, H. Farley, J. Twanmoh, A. Urumov, N. Olsen, B. Evans, and H. Kabiri. Emergency department patient flow: The influence of hospital census variables on emergency department length of stay. *Academic Emergency Medicine*, 16(7):597–602, 2009.

[37] S. Ma, M. R. Kosorok, and J. P. Fine. Additive risk models for survival data with high-dimensional covariates. *Biometrics*, 62:202–210, 2006.

[38] T. Martinussen and T. H. Scheike. Covariate selection for the semiparametric additive risk model. *Scandinavian Journal of Statistics*, 36:602–619, 2009.

[39] L. Mason, J. Baxter, P. L. Bartlett, and M. Frean. Functional gradient techniques for combining hypotheses. In *Advances in Large Margin Classifiers*. MIT Press, 2000.

[40] M. L. McCarthy, S. L. Zeger, R. Ding, S. R. Levin, J. S. Desmond, J. Lee, and D. Aronsky. Crowding delays treatment and lengthens emergency department length of stay, even among high-acuity patients. *Annals of Emergency Medicine*, 54(4):492–503, 2009.

[41] I. W. McKeague and K. J. Utikal. Inference for a nonlinear counting process regression model. *Annals of Statistics*, 18(3):1172–1187, 1990.

[42] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer, 2004.

[43] J. P. Nielsen and O. B. Linton. Kernel estimation in a nonparametric marker dependent hazard model. *Annals of Statistics*, 23(5):1735–1748, 1995.

[44] F. O'Sullivan. Fast computation of fully automated log-density and log-hazard estimators. *SIAM Journal of Scientific Statistical Computing*, 9(2): 363–379, 1988.

[45] M. L. G. Perez, L. Janys, M. D. Martinez-Miranda, and J. P. Nielsen. Bandwidth selection in marker dependent kernel hazard estimation. *Computational Statistics and Data Analysis*, 68:155–169, 2013.

[46] G. Ridgeway. The state of boosting. *Computing Science and Statistics*, 31: 172–181, 1999.

[47] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.

[48] R. E. Schapire and Y. Freund. *Boosting: Foundations and Algorithms*. MIT Press, 2014.

[49] M. Schmid and T. Hothorn. Flexible boosting of accelerated failure time models. *BMC Bioinformatics*, 9:269, 2008.

[50] A. Senthilselvan. Nonparametric estimation of hazard function from left truncated and right censored data. *Journal of Nonparametric Statistics*, 2: 29–35, 1992.

[51] L. Spierdijk. Nonparametric conditional hazard rate estimation: A local linear approach. *Computational Statistics and Data Analysis*, 52:2419–2434, 2008.

[52] M. Telgarsky. Margins, shrinkage, and boosting. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

[53] R.J. Tibshirani. Degrees of freedom and model search. *Statistica Sinica*, pages 1265–1296, 2015.

[54] R.J. Tibshirani and J. Taylor. Degrees of freedom in lasso problems. *Annals of Statistics*, pages 1198–1232, 2012.

[55] A. W. van der Vaart and J. A. Wellner. *Weak convergence and empirical processes with applications to statistics*. Springer NY, 1996.

[56] I. van Keilegom and N. Veraverbeke. Hazard rate estimation in nonparametric regression with censored data. *Ann. Inst. Statist. Math*, 53(4):730–745, 2001.

[57] Q. H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57(2):307–333, 1989.

[58] T. Zhang and B. Yu. Boosting with early stopping: Convergence and consistency. *Annals of Statistics*, 33:1538–1579, 2005.

[59] H. Zou, T.J. Hastie, and R. Tibshirani. On the 'degrees of freedom' of the lasso. *Annals of Statistics*, 35(5):2173–2192, 2007.